Sparsity-Exploiting Moment-Based Relaxations of the Optimal Power Flow Problem

Daniel K. Molzahn, *Member, IEEE* and Ian A. Hiskens, *Fellow, IEEE*

arXiv:1404.5071v2 [math.OC] 17 Nov 2014

***Abstract*—Convex relaxations of non-convex optimal power flow (OPF) problems have recently attracted significant interest. While existing relaxations globally solve many OPF problems, there are practical problems for which existing relaxations fail to yield physically meaningful solutions. This paper applies moment relaxations to solve many of these OPF problems. The moment relaxations are developed from the Lasserre hierarchy for solving generalized moment problems. Increasing the relaxation order in this hierarchy results in "tighter" relaxations at the computational cost of larger semidefinite programs. Low-order moment relaxations are capable of globally solving many small OPF problems for which existing relaxations fail. By exploiting sparsity and only applying the higher-order relaxation to specific buses, global solutions to larger problems are computationally tractable through the use of an iterative algorithm informed by a heuristic for choosing where to apply the higher-order constraints. With standard semidefinite programming solvers, the algorithm globally solves many test systems with up to 300 buses for which the existing semidefinite relaxation fails to yield globally optimal solutions.**

***Index Terms*—Optimal power flow, Semidefinite optimization, Moment relaxations, Global solution**

I. Introduction

THE optimal power flow (OPF) problem determines an optimal operating point for an electric power system in terms of a specified objective function (typically generation cost per unit time), subject to both network equality constraints (i.e., the power flow equations, which model the relationship between voltages and power injections) and engineering limits (e.g., inequality constraints on voltage magnitudes, active and reactive power generations, and line flows). While the OPF problem is often augmented with security constraints that ensure robustness to contingencies (see, e.g., [1]–[4]), the formulation considered in this research does not consider contingencies.

The OPF problem is generally non-convex due to the non-linear power flow equations [5] and may have local solutions [6]. Non-convexity of the OPF problem has made solution techniques an ongoing research topic. Many OPF solution techniques have been proposed, including successive quadratic programs, Lagrangian relaxation, genetic algorithms, particle swarm optimization, and interior point methods [7]–[12]. Some of these techniques are quite mature and capable of finding at-least-locally optimal solutions to many large-scale OPF problems with reasonable computational burden. (For instance, [3] and [4] report computationally tractable methods for finding at-least-locally-optimal solutions to security-constrained OPF problems with approximately 3,000 and 9,000 buses, respectively.)

However, while typical local solution techniques often in fact find globally optimal solutions [3], [13], they may fail to converge or converge to a local optimum. See, for instance, the example problems and discussion in [6] as well as the five-bus system in [5]. For many of these problems, MATPOWER's [10] default interior point solver with default options and a variety of initialization heuristics either fails to converge or finds locally optimal solutions. See also [14], which reports on a study by the Federal Energy Regulatory Commission of convergence characteristics for a variety of commercial solvers, OPF problem formulations, and initialization heuristics.

Recently, significant attention has focused on a semidefinite relaxation of the OPF problem [15]. If the relaxed problem satisfies a rank condition (i.e., the relaxation is said to be "exact" or "tight"), the global solution to the original OPF problem can be determined in polynomial time. Prior OPF solution methods do not guarantee finding a global solution in polynomial time. Further, infeasibility of a relaxation certifies infeasibility of the OPF problem, which is a capability not available with typical existing solution techniques. Additionally, unlike local solution techniques whose convergence characteristics generally depend on a chosen initialization, the semidefinite relaxation provides the global solution regardless of the choice of initialization when the relaxation is exact. While not as mature as existing solution techniques, semidefinite programming approaches thus have substantial advantages over traditional techniques.

However, the rank condition is not satisfied for all practical OPF problems [6], [16]. For such problems, the relaxation provides a lower bound on the optimal objective value but does not provide physically meaningful decision variables (i.e., voltage phasors). The bounds obtained from the semidefinite relaxation are often close to the global optimum and are therefore useful for many applications (e.g., calculating voltage stability margins [17] and determining the potential suboptimality of an OPF solution that is only guaranteed to be locally optimal). However, determining both the globally optimal objective value and the globally optimal decision variables is important in many contexts.

To address problems for which the rank condition is not satisfied, this paper presents *moment relaxations*[1] that globally solve a broader class of OPF problems than existing relaxations. Building on the results of [20] (many of which are independently studied in [21] and [22]), this paper describes a

University of Michigan Department of Electrical Engineering and Computer Science: molzahn@umich.edu, hiskens@umich.edu

[1]The terminology *moment relaxation*, adopted from [18], [19], refers to the relaxation's derivation from a special case of the *generalized moment problem*.

method for globally solving large OPF problems by exploiting sparsity and only applying computationally intensive "higher-order" moment relaxations to specific buses.

Much of the related work in this area focuses on sufficient conditions for which existing convex relaxations are exact [23]–[25]. While the sufficient conditions developed thus far are promising, they only apply to a limited subset of problems.

For more general cases, [26] proposes a heuristic method for finding a global optimum that is "hidden" in a higher-rank subspace of solutions to the semidefinite relaxation. However, the semidefinite relaxation yields an optimal objective value strictly less than the global minimum of some OPF problems [16]. For such cases, other heuristics may obtain at-least-locally optimal solutions [25], [26] with the optimal objective value of the semidefinite relaxation indicating the potential suboptimality.

While deserving of further study, heuristics eliminate the global optimality guarantee that is one of the main advantages of the semidefinite relaxation. We propose an alternative moment relaxation that, when exact, yields the global optimum. Using polynomial optimization theory [18], [19], moment relaxations globally solve a broad class of OPF problems, including many problems for which existing relaxations are not exact. Moment relaxations exploit the fact that the OPF problem is composed of polynomials in the voltage phasor components and is therefore a polynomial optimization problem.

Global solution of a broader class of OPF problems has a computational cost. Whereas the matrix in the semidefinite relaxation of [15] is composed of all second-order combinations of the voltage phasor components, the moment relaxation's matrices are composed of higher-degree combinations. The semidefinite program for the order-γ moment relaxation of an n-bus system has a positive semidefinite constraint on a $k \times k$ matrix, where $k = (2n+\gamma)!/\left((2n)!\gamma!\right)$ (i.e., this matrix is composed of all combinations of voltage components up to order 2γ). For example, the moment matrices for the first-, second-, and third-order moment relaxations of a 10-bus system have size 21×21, 231×231, and 1771×1771, respectively, as compared to 20×20 for the semidefinite relaxation of [15]. Thus, the computational requirements of the moment relaxations can be substantially larger than the semidefinite relaxation of [15], especially for high-order relaxations.

Fortunately, experience with small systems suggests that low-order relaxations globally solve a broad class of OPF problems, including problems for which the semidefinite relaxation of [15] is not exact. As an example of the effectiveness of the moment relaxations, consider the 9-bus OPF problem in [6]. MATPOWER [10] with the default interior point solver and default solver options either fails to converge or converges to one of three local optima depending on the initialization.[2] (The local optima have objective values that are 10.0%, 37.5%, and 38.1% greater than the global optimum.) The semidefinite relaxation of [15] yields a lower bound that is 11% less than the global optimum. Thus, an existing convex relaxation and a typical interior point technique both perform poorly for this problem while a second-order relaxation finds the global solution [20]. The capabilities of low-order relaxations for small OPF problems are further described in [20], which includes an exploration of the feasible spaces of second-order relaxations, and independently in [21] and [22].

However, large OPF problems are computationally intractable even for low-order relaxations. Solving existing semidefinite relaxations of large OPF problems requires exploiting power system sparsity. Using a matrix completion decomposition, existing semidefinite relaxations are computationally tractable for problems with thousands of buses [27], [28]. Naïve application of related techniques for the moment relaxations [29] enables solution of systems with up to approximately forty buses. Solving larger systems requires more judicious use of the higher-order relaxation. Exploiting the observation that power injection "mismatches" typically occur only in small regions of realistic large OPF problems [16], this paper applies a higher-order relaxation to specific buses. This enables global solution of large problems. With standard semidefinite programming solvers, the proposed approach is successfully applied to OPF problems with up to 300 buses for which the semidefinite relaxation of [15] fails to yield the globally optimal decision values. Further improvements in solving larger OPF problems may be achieved by combining emerging semidefinite programming solvers with the method proposed in this paper.

Selective application of the moment relaxation is independently proposed in [22]. The method described in [22] is limited to second-order relaxations of OPF problems with less than 40 buses due to a computationally expensive subproblem and lack of concurrent exploitation of sparsity.

After introducing the OPF problem formulation in Section II, we describe the moment relaxations in Section III. Section IV then presents the method for globally solving large OPF problems by exploiting sparsity and only applying the higher-order relaxations to specific buses. Section V presents results from the proposed method. Section VI concludes the paper and discusses future research directions.

II. OPF Problem Formulation

We first present an OPF formulation in terms of rectangular voltage coordinates, active and reactive power generation, and apparent-power line-flow limits. Consider an n-bus power system, where $\mathcal{N} = \{1, 2, \ldots, n\}$ is the set of all buses, $\mathcal{G}$ is the set of generator buses, and $\mathcal{L}$ is the set of all lines. Let $P_{Dk} + jQ_{Dk}$ represent the active and reactive load demand at each bus $k \in \mathcal{N}$. Let $V_k = V_{dk} + jV_{qk}$ denote the voltage phasors in rectangular coordinates at each bus $k \in \mathcal{N}$. Superscripts "max" and "min" denote specified upper and lower limits. Buses without generators have maximum and minimum generation set to zero. Let $\mathbf{Y} = \mathbf{G} + j\mathbf{B}$ denote the network admittance matrix. Shunt conductances and susceptances at bus k contribute to the diagonal element $\mathbf{Y}_{kk}$.

[2] Initialization heuristics included 1.) a "flat start" with unity voltage magnitudes and zero voltage angles, 2.) the solution to the linear "DC" OPF approximation, 3.) a power flow solution calculated using active power injections at the midpoints of the generators' operating ranges, 4.) a power flow solution calculated using power injections corresponding to an economic dispatch, and 5.) a power flow solution calculated using the power injections resulting from a DC OPF.

The power flow equations describe the network physics:

$$P_{Gk} = f_{Pk}(V_d, V_q) = V_{dk} \sum_{i=1}^{n} (\mathbf{G}_{ik} V_{di} - \mathbf{B}_{ik} V_{qi}) + V_{qk} \sum_{i=1}^{n} (\mathbf{B}_{ik} V_{di} + \mathbf{G}_{ik} V_{qi}) + P_{Dk} \quad (1a)$$

$$Q_{Gk} = f_{Qk}(V_d, V_q) = V_{dk} \sum_{i=1}^{n} (-\mathbf{B}_{ik} V_{di} - \mathbf{G}_{ik} V_{qi}) + V_{qk} \sum_{i=1}^{n} (\mathbf{G}_{ik} V_{di} - \mathbf{B}_{ik} V_{qi}) + Q_{Dk} \quad (1b)$$

Define a convex quadratic cost of active power generation:

$$f_{Ck}(V_d, V_q) = c_{k2} \left(f_{Pk}(V_d, V_q)\right)^2 + c_{k1} f_{Pk}(V_d, V_q) + c_{k0} \quad (2)$$

Note that while we focus on minimization of a quadratic function of active power generation, one can substitute other cost functions (e.g., loss minimization, voltage regulation, convex piecewise-linear generation cost functions, reactive power dispatch, etc.) for (2). (The moment relaxation approach described in this paper is applicable for any polynomial or convex piecewise-polynomial objective function.)

Define a function for squared voltage magnitude:

$$(V_k)^2 = f_{Vk}(V_d, V_q) = V_{dk}^2 + V_{qk}^2 \quad (3)$$

Squared apparent-power line flows are polynomial functions of the voltage components V_d and V_q. To account for flow limits on transformers with non-zero phase shifts and/or off-nominal voltage ratios, we model the line from bus l to bus m as a Π-model circuit with series admittance $g_{lm} + jb_{lm}$ and total shunt admittance $g_{sh,lm} + jb_{sh,lm}$ in series with an ideal transformer with a specified complex turns ratio $1 : \tau_{lm} e^{j\theta_{lm}}$ as in [10]. (Note that the conductance $g_{sh,lm}$ in the Π-model is generally neglected in typical power system data sets, and that shunt susceptances $b_{sh,lm}$ are often neglected for transformers.)

$$P_{lm} = f_{Plm}(V_d, V_q) = \left(V_{dl}^2 + V_{ql}^2\right)\left(g_{lm} + \frac{g_{sh,lm}}{2}\right)/\tau_{lm}^2 + (V_{dl}V_{dm} + V_{ql}V_{qm})(b_{lm}\sin(\theta_{lm}) - g_{lm}\cos(\theta_{lm}))/\tau_{lm} + (V_{dl}V_{qm} - V_{ql}V_{dm})(g_{lm}\sin(\theta_{lm}) + b_{lm}\cos(\theta_{lm}))/\tau_{lm} \quad (4a)$$

$$P_{ml} = f_{Pml}(V_d, V_q) = \left(V_{dm}^2 + V_{qm}^2\right)\left(g_{lm} + \frac{g_{sh,lm}}{2}\right) - (V_{dl}V_{dm} + V_{ql}V_{qm})(g_{lm}\cos(\theta_{lm}) + b_{lm}\sin(\theta_{lm}))/\tau_{lm} + (V_{dl}V_{qm} - V_{ql}V_{dm})(g_{lm}\sin(\theta_{lm}) - b_{lm}\cos(\theta_{lm}))/\tau_{lm} \quad (4b)$$

$$Q_{lm} = f_{Qlm}(V_d, V_q) = -\left(V_{dl}^2 + V_{ql}^2\right)\left(b_{lm} + \frac{b_{sh,lm}}{2}\right)/\tau_{lm}^2 + (V_{dl}V_{dm} + V_{ql}V_{qm})(b_{lm}\cos(\theta_{lm}) + g_{lm}\sin(\theta_{lm}))/\tau_{lm} + (V_{dl}V_{qm} - V_{ql}V_{dm})(g_{lm}\cos(\theta_{lm}) - b_{lm}\sin(\theta_{lm}))/\tau_{lm} \quad (4c)$$

$$Q_{ml} = f_{Qml}(V_d, V_q) = -\left(V_{dm}^2 + V_{qm}^2\right)\left(b_{lm} + \frac{b_{sh,lm}}{2}\right) + (V_{dl}V_{dm} + V_{ql}V_{qm})(b_{lm}\cos(\theta_{lm}) - g_{lm}\sin(\theta_{lm}))/\tau_{lm} + (-V_{dl}V_{qm} + V_{ql}V_{dm})(g_{lm}\cos(\theta_{lm}) + b_{lm}\sin(\theta_{lm}))/\tau_{lm} \quad (4d)$$

$$(S_{lm})^2 = f_{Slm}(V_d, V_q) = \left(f_{Plm}(V_d, V_q)\right)^2 + \left(f_{Qlm}(V_d, V_q)\right)^2 \quad (4e)$$

$$(S_{ml})^2 = f_{Sml}(V_d, V_q) = \left(f_{Pml}(V_d, V_q)\right)^2 + \left(f_{Qml}(V_d, V_q)\right)^2 \quad (4f)$$

The classical OPF problem is then

$$\min_{V_d, V_q} \sum_{k \in \mathcal{G}} f_{Ck}(V_d, V_q) \qquad \text{subject to} \quad (5a)$$

$$P_{Gk}^{\min} \le f_{Pk}(V_d, V_q) \le P_{Gk}^{\max} \qquad \forall k \in \mathcal{N} \quad (5b)$$

$$Q_{Gk}^{\min} \le f_{Qk}(V_d, V_q) \le Q_{Gk}^{\max} \qquad \forall k \in \mathcal{N} \quad (5c)$$

$$\left(V_k^{\min}\right)^2 \le f_{Vk}(V_d, V_q) \le \left(V_k^{\max}\right)^2 \qquad \forall k \in \mathcal{N} \quad (5d)$$

$$f_{Slm}(V_d, V_q) \le (S_{lm}^{\max})^2 \qquad \forall (l,m) \in \mathcal{L} \quad (5e)$$

$$f_{Sml}(V_d, V_q) \le (S_{lm}^{\max})^2 \qquad \forall (l,m) \in \mathcal{L} \quad (5f)$$

$$V_{q1} = 0 \quad (5g)$$

Constraints (5e) and (5f) limit the apparent-power flow at each line terminal. Constraint (5g) sets the reference bus angle to zero.

Note that the OPF problem is often extended to consider contingency, voltage-stability, and transient stability constraints. (See [1]–[4] for discussions of these and other extensions.) As a starting point for the methods developed in this paper, we only consider OPF problems without these constraints. Future work includes extension to more general OPF problem formulations. (See [30] and [31] for initial research on applications of convex relaxations to OPF problems with contingency constraints.)

Further, the OPF formulation studied in this paper does not consider the decision variables associated with controllable power system devices such as high-voltage DC (HVDC) lines, tap-changing and phase-shifting transformers, and switched-shunt devices. Incorporating the continuous and potentially discrete decision variables necessary for modeling these devices in convex relaxations of the OPF problem is an area of ongoing research. For instance, a convex, second-order cone programming formulation for HVDC lines is available in [32]. One possibility is modeling discrete variables as polynomial equality constraints (e.g., $\phi \in \{0, 1\}$ is equivalent to $\phi^2 - \phi = 0$), which can be incorporated in the moment relaxations discussed in this paper. This is a promising direction for future work.

III. Moment Relaxations

A. Overview

The OPF problem (5) is comprised of polynomial functions of the voltage components V_d and V_q and can therefore be solved using moment relaxations [18], [19]. We next present moment relaxations of the OPF problem (5). The material in this section builds on [20]. More detailed descriptions of moment relaxations are available in [18] and [19], and

application of moment relaxations to the OPF problem is independently proposed in [21] and [22].

Polynomial optimization problems, such as the OPF problem, are a special case of *generalized moment problems* [19]. Global solutions to generalized moment problems can be approximated using moment relaxations that are formulated as semidefinite programs. For polynomial optimization problems with bounded variables, such as OPF problems, the approximation approaches the global solution(s) as the relaxation order increases [19]. While moment relaxations can find all global solutions to polynomial optimization problems, we focus on problems with a single global optimum.

Formulating the moment relaxations requires several definitions. Define the vector $\hat{x} = \begin{bmatrix} V_{d1} & V_{d2} & \dots & V_{qn} \end{bmatrix}^\intercal$, which contains all first-order monomials of the decision variables in (5). Given a vector $\alpha \in \mathbb{N}^{2n}$ representing monomial exponents, the expression $\hat{x}^\alpha = V_{d1}^{\alpha_1} V_{d2}^{\alpha_2} \cdots V_{qn}^{\alpha_{2n}}$ defines the monomial associated with $\hat{x}$ and α. A polynomial $g(\hat{x})$ can be expressed as

$$g(\hat{x}) \triangleq \sum_{\alpha \in \mathbb{N}^{2n}} g_\alpha \hat{x}^\alpha \tag{6}$$

where g_α is the scalar coefficient corresponding to the monomial $\hat{x}^\alpha$.

Next define a linear functional $L_y\{g\}$:

$$L_y\{g\} \triangleq \sum_{\alpha \in \mathbb{N}^{2n}} g_\alpha y_\alpha \tag{7}$$

This functional replaces the monomials $\hat{x}^\alpha$ in a polynomial function $g(\hat{x})$ with scalar variables y_α. When $g(\hat{x})$ is a matrix, the functional $L_y\{g\}$ is applied to each element of $g(\hat{x})$.

Consider, for example, the vector $\hat{x} = \begin{bmatrix} V_{d1} & V_{d2} & V_{q2} \end{bmatrix}^\intercal$ corresponding to the voltage components of a two-bus system, where the angle reference constraint (5g) is used to eliminate V_{q1}, and the polynomial $g(\hat{x}) = -(0.95)^2 + f_{V2}(V_d, V_q) = -(0.95)^2 + V_{d2}^2 + V_{q2}^2$. (The constraint $g(\hat{x}) \geq 0$ forces the voltage magnitude at bus 2 to be greater than or equal to 0.95 per unit.) Then $L_y\{g\} = -(0.95)^2 y_{000} + y_{020} + y_{002}$. Thus, $L\{g\}$ converts a polynomial $g(\hat{x})$ to a linear function of y.

The order-γ relaxation forms a vector x_γ composed of all monomials of the voltage components up to order γ:

$$x_\gamma \triangleq \begin{bmatrix} 1 & V_{d1} & \dots & V_{qn} & V_{d1}^2 & V_{d1}V_{d2} & \dots \\ & \dots & V_{qn}^2 & V_{d1}^3 & V_{d1}^2 V_{d2} & \dots & V_{qn}^\gamma \end{bmatrix}^\intercal \tag{8}$$

We now define *moment* and *localizing* matrices. The symmetric moment matrix $\mathbf{M}_\gamma(y)$ has entries y_α corresponding to all monomials $\hat{x}^\alpha$ up to order 2γ:

$$\mathbf{M}_\gamma\{y\} \triangleq L_y\left\{x_\gamma x_\gamma^\intercal\right\} \tag{9}$$

Symmetric localizing matrices[3] are defined for each constraint of (5). The localizing matrices consist of linear combinations of the moment matrix entries y. Each polynomial constraint of the form $f(\hat{x}) - a \geq 0$ in (5) (e.g., $f_{V2}(\hat{x}) - V_2^{\min} \geq 0$) corresponds to the localizing matrix

$$\mathbf{M}_{\gamma-\beta}\{(f(\hat{x}) - a)\, y\} \triangleq L_y\left\{(f(\hat{x}) - a)\, x_{\gamma-\beta} x_{\gamma-\beta}^\intercal\right\} \tag{10}$$

where the polynomial f has degree 2β. Example moment and localizing matrices for the second-order relaxation of a two-bus system are presented in (13) and (14), respectively.

The order-γ moment relaxation of (5) is

$$\min_y L_y\left\{\sum_{k \in \mathcal{G}} f_{Ck}\right\} \quad \text{subject to} \tag{11a}$$
$$\mathbf{M}_{\gamma-1}\left\{\left(f_{Pk} - P_k^{\min}\right) y\right\} \succeq 0 \qquad \forall k \in \mathcal{N} \tag{11b}$$
$$\mathbf{M}_{\gamma-1}\left\{\left(P_k^{\max} - f_{Pk}\right) y\right\} \succeq 0 \qquad \forall k \in \mathcal{N} \tag{11c}$$
$$\mathbf{M}_{\gamma-1}\left\{\left(f_{Qk} - Q_k^{\min}\right) y\right\} \succeq 0 \qquad \forall k \in \mathcal{N} \tag{11d}$$
$$\mathbf{M}_{\gamma-1}\left\{\left(Q_k^{\max} - f_{Qk}\right) y\right\} \succeq 0 \qquad \forall k \in \mathcal{N} \tag{11e}$$
$$\mathbf{M}_{\gamma-1}\left\{\left(f_{Vk} - V_k^{\min}\right) y\right\} \succeq 0 \qquad \forall k \in \mathcal{N} \tag{11f}$$
$$\mathbf{M}_{\gamma-1}\left\{\left(V_k^{\max} - f_{Vk}\right) y\right\} \succeq 0 \qquad \forall k \in \mathcal{N} \tag{11g}$$
$$\mathbf{M}_{\gamma-2}\left\{\left(S_{lm}^{\max} - f_{Slm}\right) y\right\} \succeq 0 \qquad \forall (l,m) \in \mathcal{L} \tag{11h}$$
$$\mathbf{M}_{\gamma-2}\left\{\left(S_{lm}^{\max} - f_{Sml}\right) y\right\} \succeq 0 \qquad \forall (l,m) \in \mathcal{L} \tag{11i}$$
$$\mathbf{M}_\gamma(y) \succeq 0 \tag{11j}$$
$$y_{00\dots0} = 1 \tag{11k}$$
$$y_{0\dots00\eta0\dots0} = 0 \qquad \eta = 1, \dots, 2\gamma \tag{11l}$$

where $\succeq 0$ indicates that the corresponding matrix is positive semidefinite. The moment relaxation is thus a semidefinite program. (A dual form of the moment relaxation is a sum-of-squares program [19].) Note that the constraint (11k) enforces the fact that $x^0 = 1$. The constraint (11l) corresponds to the angle reference constraint (5g); the η in (11l) is in the index $n+1$, which corresponds to the variable V_{q1}. Note that the angle reference can alternatively be used to eliminate all terms corresponding to V_{q1} to reduce the size of the semidefinite program.

B. Two-Bus Example

We next present an illustrative two-bus example problem from [33]. Fig. 1 gives the system's one-line diagram assuming a 100 MVA base power. The generator at bus 1 has no limits on active or reactive outputs and there is no line-flow limit. Bus 1 voltage magnitude is in the range $[0.95, 1.05]$ per unit, while bus 2 voltage magnitude is in the range $[0.95, 1.02]$ per unit. Specify a \$1/MWh cost of active power generation at bus 1. With three degrees of freedom, the entire feasible space of the

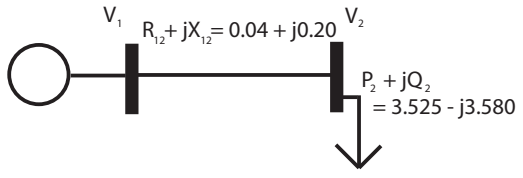

Fig. 1. Two-Bus System from [33]

[3]The terminology *localizing matrix* is adopted from [18], [19].

$$x_2 = \begin{bmatrix} 1 & V_{d1} & V_{d2} & V_{q2} & V_{d1}^2 & V_{d1}V_{d2} & V_{d1}V_{q2} & V_{d2}^2 & V_{d2}V_{q2} & V_{q2}^2 \end{bmatrix}^\intercal \tag{12}$$

$$\mathbf{M}_2(y) = L_y(x_2 x_2^\intercal) = \left[\begin{array}{c|c|c|c|c|c|c|c|c|c} y_{000} & y_{100} & y_{010} & y_{001} & y_{200} & y_{110} & y_{101} & y_{020} & y_{011} & y_{002} \\ \hline y_{100} & y_{200} & y_{110} & y_{101} & y_{300} & y_{210} & y_{201} & y_{120} & y_{111} & y_{102} \\ y_{010} & y_{110} & y_{020} & y_{011} & y_{210} & y_{120} & y_{111} & y_{030} & y_{021} & y_{012} \\ y_{001} & y_{101} & y_{011} & y_{002} & y_{201} & y_{111} & y_{102} & y_{021} & y_{012} & y_{003} \\ \hline y_{200} & y_{300} & y_{210} & y_{201} & y_{400} & y_{310} & y_{301} & y_{220} & y_{211} & y_{202} \\ y_{110} & y_{210} & y_{120} & y_{111} & y_{310} & y_{220} & y_{211} & y_{130} & y_{121} & y_{112} \\ y_{101} & y_{201} & y_{111} & y_{102} & y_{301} & y_{211} & y_{202} & y_{121} & y_{112} & y_{103} \\ y_{020} & y_{120} & y_{030} & y_{021} & y_{220} & y_{130} & y_{121} & y_{040} & y_{031} & y_{022} \\ y_{011} & y_{111} & y_{021} & y_{012} & y_{211} & y_{121} & y_{112} & y_{031} & y_{022} & y_{013} \\ y_{002} & y_{102} & y_{012} & y_{003} & y_{202} & y_{112} & y_{103} & y_{022} & y_{013} & y_{004} \end{array}\right] \tag{13}$$

$$\mathbf{M}_1\left\{\left(f_{V2} - (0.95)^2\right) y\right\} =$$

$$\left[\begin{array}{c|c|c} y_{020} + y_{002} - (0.95)^2 y_{000} & y_{120} + y_{102} - (0.95)^2 y_{100} & y_{030} + y_{012} - (0.95)^2 y_{010} \quad y_{021} + y_{003} - (0.95)^2 y_{001} \\ \hline y_{120} + y_{102} - (0.95)^2 y_{100} & y_{220} + y_{202} - (0.95)^2 y_{200} & y_{130} + y_{112} - (0.95)^2 y_{110} \quad y_{121} + y_{103} - (0.95)^2 y_{101} \\ y_{030} + y_{012} - (0.95)^2 y_{010} & y_{130} + y_{112} - (0.95)^2 y_{110} & y_{040} + y_{022} - (0.95)^2 y_{020} \quad y_{031} + y_{013} - (0.95)^2 y_{011} \\ y_{021} + y_{003} - (0.95)^2 y_{001} & y_{121} + y_{103} - (0.95)^2 y_{101} & y_{031} + y_{013} - (0.95)^2 y_{011} \quad y_{022} + y_{004} - (0.95)^2 y_{002} \end{array}\right] \tag{14}$$

two-bus OPF problem can be visualized in three dimensions as shown in Section IV of [20].

Note that while this system has a radial network topology, the OPF problem does not satisfy the sufficient conditions for exactness of the semidefinite relaxation described in [23], [24], and the first-order moment relaxation is, in fact, not exact for this problem.

The second-order relaxation has the vector x_2 given in (12) and the moment matrix in (13). The localizing matrix corresponding to the voltage magnitude constraint $f_{V2}(V_d, V_q) \geq (0.95)^2$ in (5d) is given in (14). Note that the angle reference constraint (5g) is used to eliminate V_{q1} so that $\hat{x} = \begin{bmatrix} V_{d1} & V_{d2} & V_{q2} \end{bmatrix}^\intercal$.

The second-order moment relaxation yields the global solution of $V = \begin{bmatrix} 0.950 & 0.416 - j0.893 \end{bmatrix}^\intercal$ per unit. This corresponds to active and reactive power generation at bus 1 of 456.6 MW and 162.3 MVAr, respectively, and an operating cost of 456.55 \$/MWh.

C. Implementation Details

This section next discusses several implementation details for the moment relaxations, including removing linear dependency resulting from equality-constrained polynomials, recovering the global solution(s) to the OPF problem from a solution to the moment relaxation, elimination of unnecessary portions of the moment relaxation's matrices, and reformulation of the cost function and apparent-power line-flow constraints. Note that, in contrast to the example in Section III-B, the notation in this section does not use the angle reference constraint $V_{q1} = 0$ to eliminate V_{q1}.

1) Equality-Constrained Polynomials: Given that the localizing matrices are symmetric, the constraints $\mathbf{M}_{\gamma-1}\{(f(\hat{x}) - a)\, y\} \succeq 0$ and $\mathbf{M}_{\gamma-1}\{(a - f(\hat{x}))\, y\} \succeq 0$ imply that the matrix $\mathbf{M}_{\gamma-1}\{(f(\hat{x}) - a)\, y\} = 0$. Thus, all entries of a localizing matrix corresponding to an equality-constrained polynomial $f(\hat{x}) = a$ (e.g., the power flow constraints at load buses) are zero. That is, equality-constrained polynomials result in equality constraints that are linear in the variables y rather than positive semidefinite matrix constraints.

The non-uniqueness of the entries of the localizing matrices (see, e.g., the repeated terms in (14) due to the matrix's symmetry) creates linearly dependent equality constraints which can introduce numerical difficulties. To eliminate these redundant constraints, the localizing matrix constraints for equality-constrained polynomials are therefore replaced by the equivalent vector of constraints $L_y\{g(\hat{x})\, x_{\gamma-1}\} = 0$, where $g(\hat{x})$ denotes any equality-constrained polynomial in (5).

2) Solution Extraction: The order-γ moment relaxation yields a single global solution if $\operatorname{rank}(\mathbf{M}_\gamma(y)) = 1$. The global solution x^* to the OPF problem (5) is then determined by a spectral decomposition of the diagonal block of the moment matrix corresponding to the second-order terms. Specifically, let η be a unit-length eigenvector corresponding to the non-zero eigenvalue λ from the diagonal block of the moment matrix corresponding to the second-order monomials (i.e., $[\mathbf{M}_2]_{(2:k,2:k)}$, where $k = 2n + 1$ and subscripts indicate the vector entries in MATLAB notation). Then the vector $V^* = \sqrt{\lambda}\left(\eta_{1:n} + j\eta_{(n+1):2n}\right)$ is the globally optimal voltage phasor vector.

If $\operatorname{rank}(\mathbf{M}_\gamma(y)) > 1$, there are either multiple global

solutions (i.e., there are multiple points in the feasible space of the original non-convex OPF problem with the same globally optimal objective value) requiring the solution extraction procedure described in Section 5.3.1 of [19][4] or the order-γ moment relaxation is not exact and only yields a lower bound on the objective value. If the order-γ moment relaxation is not exact, the order-$(\gamma+1)$ moment relaxation will improve the lower bound and may give a global solution.

3) Elimination of Unnecessary Terms: Since the polynomials in (5) are composed solely of constant, second-, and fourth-order monomials, off-diagonal blocks of the moment matrix corresponding to odd-order monomials are not required. Further, all terms in the off-diagonal blocks of the moment matrix corresponding to even-order monomials are duplicated in the diagonal blocks and are therefore unnecessary. (See (13) for an illustration of this matrix partitioning.) Thus, positive semidefinite matrix constraints are only enforced for the diagonal blocks of the moment matrix corresponding to even-order monomials (i.e., y_α such that $\sum_{i=1}^{2n}\alpha_i$ is even). Similarly, positive semidefinite constraints are only applied to the diagonal blocks of the localizing matrices which correspond to the even-order monomials of the matrix $x_{\gamma-\beta}x_{\gamma-\beta}^\intercal$ (e.g., the diagonal blocks of (14)). By reducing the size of the semidefinite program, this decreases the relaxation's computational burden.

After elimination of terms corresponding to the first-order monomials, the moment matrix for the first-order relaxation contains only the second-order monomials. Further, the localizing "matrices" are in fact positivity constraints on scalars.[5] Thus, the first-order relaxation is equivalent to the semidefinite relaxation of [15]. The higher-order moment relaxations are generalizations of the semidefinite relaxation of [15].

4) Quadratic Cost Function and Apparent-Power Line-FLow Limits: The order γ of the moment relaxation must be greater than or equal to half of the degree of any polynomial in the OPF problem (5). Relaxations of all polynomials can then be written as linear functions of the entries of $\mathbf{M}_\gamma$. For instance, the OPF problem with a linear cost function and without apparent-power line-flow limits requires $\gamma \geq 1$. Although direct implementation of (5) requires $\gamma \geq 2$ due to the fourth-order polynomials in the cost function (5a) and apparent-power line-flow limits (5e)-(5f), these fourth-order polynomials can be rewritten as second-order polynomials using a Schur complement formulation [15]. Specifically, rather than (11h)-(11i), enforce the constraints

$$\begin{bmatrix} -\left(S_{lm}^{\max}\right)^2 & L_y\left\{f_{Plm}\right\} & L_y\left\{f_{Qlm}\right\} \\ L_y\left\{f_{Plm}\right\} & -1 & 0 \\ L_y\left\{f_{Qlm}\right\} & 0 & -1 \end{bmatrix} \preceq 0 \quad \forall\,(l,m)\in\mathcal{L} \tag{15a}$$

$$\begin{bmatrix} -\left(S_{lm}^{\max}\right)^2 & L_y\left\{f_{Pml}\right\} & L_y\left\{f_{Qml}\right\} \\ L_y\left\{f_{Pml}\right\} & -1 & 0 \\ L_y\left\{f_{Qml}\right\} & 0 & -1 \end{bmatrix} \preceq 0 \quad \forall\,(l,m)\in\mathcal{L} \tag{15b}$$

Similarly, define new variables α_k for each generator $k\in\mathcal{G}$ and replace the quadratic cost function in (11a) with $\sum_{k\in\mathcal{G}}\alpha_k$ and the additional constraint

$$\begin{bmatrix} c_{k1}L_y\left\{f_{Pk}\right\}+c_{k0}-\alpha_k & \sqrt{c_{k2}}L_y\left\{f_{Pk}\right\} \\ \sqrt{c_{k2}}L_y\left\{f_{Pk}\right\} & -1 \end{bmatrix} \preceq 0 \quad \forall k\in\mathcal{G} \tag{16}$$

Note that a second-order cone programming (SOCP) formulation can also be employed to represent the quadratic cost function and apparent-power line-flow limits [34].

The OPF problem reformulated using (15) and (16) only requires $\gamma \geq 1$. Use of both the Schur complement formulations and direct implementation for the apparent-power line-flow constraints (5e)-(5f) and quadratic cost function (11a) generally gives superior results for $\gamma \geq 2$ as compared to implementing either the Schur complement or direct formulations separately. That is, when possible, enforce both (11h)-(11i) and (15) for the apparent-power line-flow constraints and include the constraints (16) and $\alpha_k = L\left\{f_{Ck}\right\}\ \forall k\in\mathcal{G}$ while minimizing $\sum_{k\in\mathcal{G}}\alpha_k$ for the quadratic cost function).

IV. Exploiting Sparsity in Moment Relaxations

The moment relaxations globally solve a broader class of OPF problems than existing convex relaxations [20]–[22]. However, the superior capabilities of the moment relaxations have a computational cost: the semidefinite program needed to solve the moment relaxation quickly becomes computationally intractable with both increasing problem size and relaxation order. Direct implementation of the formulation presented in Section III is computationally tractable for a second-order relaxation of OPF problems with up to ten buses.

Solving larger OPF problems requires exploiting power system sparsity. Similar to methods for existing semidefinite relaxations [27], [28], matrix completion decomposition is applicable to moment relaxations of polynomial optimization problems [29]. This decomposition extends computational tractability to OPF problems with approximately forty buses.

Solving OPF problems with more than forty buses requires exploiting the observation that the first-order relaxation is sufficient for large regions of typical OPF problems [16]. By selectively applying second- and third-order relaxations to specific buses or small groups of buses, larger OPF problems become computationally tractable. With standard semidefinite programming solvers, OPF problems of up to 300 buses can be solved.

[4]If $\operatorname{rank}\left(\mathbf{M}_{\gamma-1}\left(y\right)\right) = \operatorname{rank}\left(\mathbf{M}_{\gamma}\left(y\right)\right)$, then there are at least $\operatorname{rank}\left(\mathbf{M}_{\gamma}\left(y\right)\right)$ globally optimal solutions to the OPF problem. The globally optimal decision variable vectors can be extracted using Algorithm 4.2 in [19], which only uses linear algebra operations.

While practical OPF problems can have multiple *local* solutions, we expect that multiple *global* solutions are uncommon. Reference [15] uses example OPF problems with multiple global solutions to show that the OPF problem is, in general, NP-hard. As discussed in [20], these example problems are atypical, but are useful for exploring the limits of the moment relaxation approach.

[5]For the first-order moment relaxation, the localizing matrix for the polynomial constraint $f\left(x\right)\geq 0$ is $L_y\left(f\left(x\right)x_0x_0^\intercal\right) = L_y\left(f\left(x\right)\cdot 1\cdot 1^\intercal\right) = L_y\left(f\left(x\right)\right)\geq 0$ (i.e., a positive scalar constraint).

After reviewing the matrix completion decomposition [29], this section proposes a method for selectively applying the moment relaxation and presents a heuristic method for determining where to apply higher-order relaxations.

A. Matrix Completion Decomposition

The matrix completion decomposition, which is adopted from [29], exploits power system sparsity. The decomposition relies on a matrix completion theorem [35], [36] which draws on graph theory. Several graph theoretic definitions are necessary for understanding the matrix completion theorem. A *clique* is a subset of the graph nodes for which each node in the clique is connected to all other nodes in the clique. A *maximal clique* is a clique that is not a proper subset of another clique. Denote the set of maximal cliques by $\mathcal{M}$, with $\mathcal{M}_m$ representing the set of buses associated with the m^{th} maximal clique. A graph is *chordal* if each cycle of length four or more nodes has a chord, which is an edge connecting two nodes that are not adjacent in the cycle.

The graph in question for the moment relaxations of the OPF problem is defined with a set of nodes $\hat{\mathcal{N}}$ and a set of undirected edges $\hat{\mathcal{L}}$. This graph is derived from the power system network. The set of nodes is equal to the set of buses in the power system (i.e., $\hat{\mathcal{N}} = \mathcal{N} = \{1, \ldots, n\}$). The set of edges $\hat{\mathcal{L}}$ is a superset of the topology of the power system network $\mathcal{L}$. Define $\bar{\mathcal{N}}_k$ as the subset of buses connected to bus k in the power system network (i.e., $\bar{\mathcal{N}}_k = \{i \mid (i,k) \in \mathcal{L}\}$). For each bus k, add to $\hat{\mathcal{L}}$ all edges between each bus in $\bar{\mathcal{N}}_k$. That is, all neighboring buses of each bus are connected in $\hat{\mathcal{L}}$.[6]

The maximal cliques of a chordal graph can be determined in linear time [37]. However, identifying the maximal cliques of a non-chordal graph is an NP-hard problem. Since realistic power networks are generally not chordal, we use a chordal extension technique which adds edges to $\hat{\mathcal{L}}$ to obtain a chordal super-graph denoted as $\hat{\mathcal{L}}_{ch}$. To form the chordal extension, denote as $\mathbf{D}$ the adjacency matrix of the graph defined by $\hat{\mathcal{N}}$ and $\hat{\mathcal{L}}$. The chordal extension is then determined using a Cholesky factorization of $\mathbf{D} + \mathbf{I}$, where $\mathbf{I}$ is an $n \times n$ identity matrix. The off-diagonal sparsity pattern of $\mathrm{chol}\,(\mathbf{D} + \mathbf{I})$ provides a chordal extension $\hat{\mathcal{L}}_{ch}$. An approximate minimum-degree permutation of the buses [38] is employed to reduce the number of added edges in $\hat{\mathcal{L}}_{ch}$ relative to $\hat{\mathcal{L}}$.

The matrix completion theorem can now be stated. Let $\mathbf{W}$ be a symmetric matrix with partial information (i.e., not all entries of $\mathbf{W}$ have known values) with an associated undirected graph. (For the moment relaxation, the graph in question has nodes $\hat{\mathcal{N}}$ and edges $\hat{\mathcal{L}}_{ch}$.) The matrix $\mathbf{W}$ can be completed to a positive semidefinite matrix (i.e., the unknown entries of $\mathbf{W}$ can be chosen such that $\mathbf{W} \succeq 0$) if and only if the submatrices associated with each of the maximal cliques of the graph defined by $\mathbf{W}$ are all positive semidefinite.

[6]This is a subtle but important difference from the matrix completion decompositions in [27] and [28] that directly use the graph with nodes $\mathcal{N}$ and edges $\mathcal{L}$ from the power system network. Matrix completion decompositions for the higher-order moment formulations use the graph defined by $\hat{\mathcal{N}}$ and $\hat{\mathcal{L}}$ so that each bus belongs to at least one maximal clique that also contains each of that bus' neighbors, and thus all variables necessary for the higher-order moment constraints are well-defined [29].

The matrix completion theorem allows replacing the single large positive semidefinite constraint on the moment matrix (11j) with constraints on many smaller matrices:

$$\mathbf{M}_{\gamma}^{\mathcal{M}_m}(y) \triangleq L_y\left\{\left(x_{\gamma}^{\mathcal{M}_m}\right)\left(x_{\gamma}^{\mathcal{M}_m}\right)^{\mathsf{T}}\right\} \succeq 0 \quad m = 1, \ldots, |\mathcal{M}| \tag{17}$$

where $x_{\gamma}^{\mathcal{M}_m}$ is the subset of x_{γ} corresponding to the buses in $\mathcal{M}_m$ and $|\mathcal{M}|$ is the number of maximal cliques in the graph.

Similarly, the localizing matrices in (11b)-(11i) are each replaced by a single smaller matrix. Each bus k is associated with a single *smallest covering* maximal clique $m_k \in \{1, \ldots, |\mathcal{M}|\}$ (i.e., the maximal clique with least number of buses that completely contains bus k and its neighbors in the power system network). By construction of $\hat{\mathcal{L}}$, each bus and its neighbors will be entirely contained in at least one maximal clique. Form the localizing matrix

$$\begin{aligned}\mathbf{M}_{\gamma-\beta}^{\mathcal{M}_{m_k}}&\{(f(\hat{x})-a)\,y\}\\ &\triangleq L_y\left\{(f(\hat{x})-a)\left(x_{\gamma-\beta}^{\mathcal{M}_{m_k}}\right)\left(x_{\gamma-\beta}^{\mathcal{M}_{m_k}}\right)^{\mathsf{T}}\right\}\end{aligned} \tag{18}$$

where $f(\hat{x}) - a \geq 0$ denotes a generic polynomial constraint in (5) with order 2β associated with bus k.

Since the maximal cliques have non-empty intersection (i.e., contain some of the same buses), different decomposed moment matrices may contain elements that refer to a common element in the original moment matrix. The decomposed optimization problem must be formulated such that these shared elements are equal.

The solution to the decomposed formulation consists of many matrices. The globally optimal voltage vector solution to the OPF problem can be recovered if each moment matrix satisfies a rank condition. Specifically, the diagonal blocks corresponding to the second-order monomials in all moment matrices must have rank one. If the rank condition is satisfied, the method described in [28] may be used to recover the globally optimal voltage vector.

If any of the decomposed moment matrices does not satisfy the rank condition, the decomposed moment relaxation does not yield a solution to the OPF problem. Failure to satisfy the rank condition may either indicate that the moment relaxation is not exact or that there are multiple global solutions. In the former case, the objective value from the moment relaxation serves as a lower bound on the objective value of the OPF problem (5) and increasing the relaxation order may result in a relaxation that yields a global solution.[7] In the latter case, unlike the formulation in Section III which can recover multiple global solutions using the method described in [19], there is only limited ability to recover multiple global solutions to the decomposed moment relaxation. Adding a small perturbation to the objective function may result in recovery of a single global solution [29].

[7]The lower bound may in fact be the global minimum objective value without the relaxation providing globally optimal decision variables (i.e., a "hidden" rank one solution [26]).

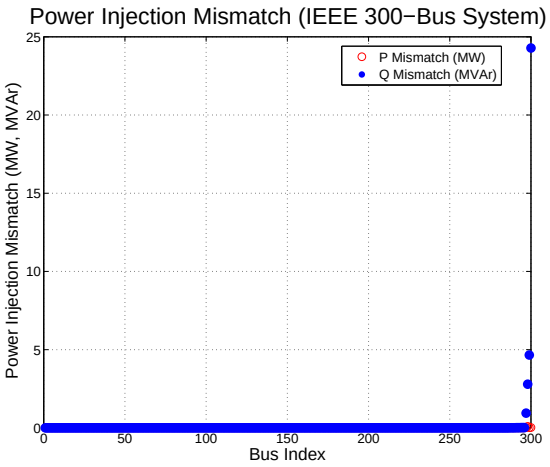

Fig. 2. Power Injection Mismatches for the First-Order Relaxation of the IEEE 300-Bus System [16]

B. Selective Application of Higher-Order Constraints

The matrix completion decomposition described in Section IV-A significantly reduces the size of the semidefinite program for large, sparse power networks. With this decomposition, second-order relaxations of OPF problems with up to approximately forty buses are computationally tractable.

Solving larger OPF problems is accomplished by exploiting the observation that the first-order relaxation is sufficient for large regions of typical OPF problems. A voltage vector is obtained from the closest matrix that satisfies the rank condition (i.e., the rank-one matrix with smallest Frobenius-norm difference to the moment matrix from the relaxation's solution). This matrix is determined using an eigen decomposition of the higher-rank diagonal block of the moment matrix corresponding to the second-order monomials. A power injection "mismatch" is determined by comparing the value of the power injections calculated from the higher-rank matrix ($L_y\{f_{Pk}\}$ and $L_y\{f_{Qk}\}$) to the power injections implied by the voltage vector from the closest matrix satisfying the rank condition.[8]

First-order relaxations typically yield voltage vectors that have small power injection mismatches at the majority of buses while a few buses have large mismatch [16]. For example, Fig. 2 shows the power injection mismatches, sorted in increasing order, resulting from the first-order moment relaxation of the IEEE 300-bus system. This suggests that selective application of higher-order relaxations to specified buses may be sufficient to globally solve larger OPF problems.

To selectively apply the higher-order constraints, each bus k has an associated relaxation order γ_k rather than a uniform order for the entire OPF problem. (A heuristic for specifying the γ_k values is presented in Section IV-C.) Determine the relaxation order for each maximal clique $\mu = 1,\ldots,|\mathcal{M}|$, denoted as $\hat{\gamma}_\mu$, which is the highest relaxation order of any bus for which μ is the smallest covering maximal clique.[9] The decomposed moment matrix constraints in (17) are formed according to this order:

$$\mathbf{M}_{\hat{\gamma}_\mu}^{\mathcal{M}_\mu}(y) \triangleq L_y\left\{\left(x_{\hat{\gamma}_\mu}^{\mathcal{M}_\mu}\right)\left(x_{\hat{\gamma}_\mu}^{\mathcal{M}_\mu}\right)^{\mathsf{T}}\right\} \succeq 0 \quad \mu = 1,\ldots,|\mathcal{M}| \tag{19}$$

The localizing matrices for the power injection and voltage magnitude constraints (11b)-(11g) at bus k are constructed according to the corresponding bus order γ_k, while the localizing matrices for the apparent-power line-flow constraints are constructed according to the highest order of either terminal bus. That is, for each constraint $f(\hat{x}) - a \geq 0$ with order 2β in (5), create the localizing matrix constraint

$$\begin{aligned}&\mathbf{M}_{\gamma_k-\beta}^{\mathcal{M}_m}\{(f(\hat{x})-a)\,y\}\\ &\quad\triangleq L_y\left\{(f(\hat{x})-a)\left(x_{\gamma_k-\beta}^{\mathcal{M}_m}\right)\left(x_{\gamma_k-\beta}^{\mathcal{M}_m}\right)^{\mathsf{T}}\right\} \succeq 0\end{aligned} \tag{20}$$

In this way, the relaxation order for the majority of buses in a large OPF problem can be set to a computationally tractable value with the computationally intensive higher-order relaxations only applied where necessary.

C. Iterative Solution Algorithm with a Heuristic for Determining the Relaxation Order

With a method for selectively applying the higher-order constraints to specific buses of an OPF problem, we next present an iterative solution algorithm for the moment relaxation. At each iteration, the algorithm uses a heuristic for specifying the value of γ_k for each bus k. Denote γ as the vector containing γ_k, $k = 1,\ldots,n$.

Algorithm 1 Iterative Solution for Moment Relaxation

1: Set $\gamma_k = 1 \quad \forall k \in \mathcal{N}$
2: **repeat**
3: Solve moment relaxation with γ
4: Calculate power injection mismatches
5: Increase the entries γ according to the heuristic
6: **until** Tolerances are satisfied
7: Calculate optimal voltage profile

Algorithm 1 iteratively solves the moment relaxation and determines the power injection mismatches. The initial relaxation order γ is set to one at every bus. If the solution meets specified tolerance criteria, the algorithm recovers the optimal voltage vector using an eigen decomposition of the diagonal block of the decomposed moment matrices corresponding to the second-order monomials. Otherwise, the relaxation order is increased at a subset of the buses with greatest power injection mismatch. Each iteration of the algorithm tightens

[8] At load buses, the power injection mismatches are equal to the difference between the specified load demands and the power injections implied by the closest rank one matrix. At generator buses, the mismatches are equal to the difference between the power injections derived from the localizing matrices (i.e., the elements in the (1,1) position of (11b) and (11d) plus the load demands) and the power injections implied by the closest rank one matrix.

[9] Note that $\hat{\gamma}_\mu$ is not necessarily the highest relaxation order of the buses in maximal clique μ. The maximal clique μ may contain buses which are also contained in other maximal cliques. The order $\hat{\gamma}_\mu$ is determined by the buses for which μ is the *smallest covering* maximal clique (i.e., μ is the smallest clique to contain that bus and all of its neighbors).

the relaxation by adding higher-order constraints. The buses with greatest mismatch typically change with the relaxation order, thus potentially requiring multiple iterations of the loop in the algorithm.

There are several tolerance criteria used to evaluate the optimality of a solution. Due to numerical inaccuracies, no solver provides a solution that exactly satisfies the rank condition. One measure of the optimality of a candidate solution is based on power injection mismatches. Let P_k^{mis} and Q_k^{mis} be the active and reactive power mismatches at bus k, respectively, resulting from the voltage vector derived from the closest matrix satisfying the rank condition. A voltage vector is accepted upon satisfaction of several convergence criteria: 1.) all apparent-power injection mismatches $S_k^{mis} = \sqrt{\left(P_k^{mis}\right)^2 + \left(Q_k^{mis}\right)^2}$ are less than a specified tolerance, 2.) the voltage magnitudes, power injections, and line flows satisfy the inequality constraints in (5) to within specified tolerances, and 3.) the objective function evaluated with the voltage vector is equal to the optimal objective function from the moment relaxation to within a specified tolerance.

We finally describe the heuristic used to update γ. Define $\gamma^{\max}$ as the highest relaxation order among all buses (i.e., $\gamma^{\max} = \max_k \gamma_k$). (Note that $\gamma^{\max}$ is not a specified limit but rather can change as the algorithm progresses.) At each iteration of the algorithm, increment γ_k at up to h buses, where h is a specified parameter, that have the largest apparent-power injection mismatches S_k^{mis} among buses satisfying two conditions: 1.) γ_k is strictly less than $\gamma^{\max}$, and 2.) S_k^{mis} is greater than the specified tolerance. If no buses satisfy these two conditions, increment γ_k at up to h buses with the largest S_k^{mis} greater than the specified tolerance and increment $\gamma^{\max}$. That is, in order to avoid unnecessarily increasing the size of the moment matrices, the heuristic avoids incrementing the maximum relaxation order $\gamma^{\max}$ until $\gamma_k = \gamma^{\max}$ at all buses k with mismatch S_k^{mis} greater than the mismatch tolerance.

There is a computational trade-off in choosing the value of h. Larger values of h likely result in fewer iterations of the algorithm but each iteration is slower if more buses than necessary have high-order relaxations. Smaller values of h result in faster solution at each iteration, but may require more iterations. Experience indicates that $h = 2$ is a good balance.

Note that this heuristic is just one of many possible approaches for specifying the relaxation order γ_k at each bus k. In addition to further analysis for varying h, future work includes comparison of this heuristic to alternative approaches.

V. Numerical Results

This section illustrates the effectiveness of the proposed algorithm for the higher-order relaxations relative to the first-order relaxation by considering several test problems. For these problems, second- and third-order relaxations are exact (i.e., the relaxation provides a solution with the globally minimal objective value and the globally optimal decision variables).

By modifying examples in the existing literature, it is straightforward to find moderate-size test problems for which existing solvers, such as the default interior point solver in MATPOWER [10], fail for a variety of reasonable initialization heuristics. For instance, in an experiment conducted by randomly perturbing the cost function in the modified 118-bus system from [6], the default interior point solver in MATPOWER with default solver options either fails to converge or converges to a local optimum in 6.9% of 10000 tested problems initialized using five typical heuristics: 1.) a "flat start" with unity voltage magnitudes and zero voltage angles, 2.) the solution to the linear "DC" OPF approximation, 3.) a power flow solution calculated using active power injections at the midpoints of the generators' operating ranges, 4.) a power flow solution calculated using power injections corresponding to an economic dispatch, and 5.) a power flow solution calculated using the power injections resulting from a DC OPF. The first-order moment relaxation succeeds in globally solving all of the modified 118-bus test problems and other test problems for which traditional solution methods fail. These test problems are not of direct interest for our purposes but provide the context for considering the higher-order moment relaxations.

Accordingly, it is also straightforward to modify examples in the existing literature to obtain test problems for which the first-order relaxation fails, but second- or third-order relaxations succeed. However, traditional solution methods, such as the interior point solver in MATPOWER, succeed in finding what turns out to be the global solution for these test problems. With the focus of this paper on demonstrating the effectiveness of higher-order moment relaxations relative to the first-order relaxation (and, equivalently, relative to the existing semidefinite relaxation [15]), we investigate these test

TABLE I
TEST CASE DESCRIPTIONS

Test Case	Description	Opt. Obj. Val. ($/hr)
case14Q	IEEE 14-bus system with all active and reactive loads decreased by 50%	3.302×10^3
case14L	IEEE 14-bus system with 25 MVA apparent-power flow limits for all lines	9.359×10^3
case39Q	39-bus system from [6], which is the IEEE 39-bus system with active and reactive loads decreased by 50% and voltage bounds tightened to $[1.05, 0.95]$ per unit	1.122×10^4
case39L	IEEE 39-bus system with all apparent-power line-flow limits decreased by 15%	4.192×10^4
case57Q	IEEE 57-bus system with active and reactive demand reduced by 75% and all lower limits on generator reactive power injections set to -10 MVAr	7.352×10^3
case57L	IEEE 57-bus system with 77 MVA apparent-power flow limits for all lines	4.398×10^4
case118Q	IEEE 118-bus system with active and reactive demand reduced by 30% and all lower limits on generator reactive power injection set to -20 MVAr	8.151×10^4
case118L	IEEE 118-bus system with 110 MVA apparent-power flow limits for all lines	1.349×10^5
case300	IEEE 300-bus system	7.200×10^5
Small minimum resistances of 1×10^{-4} per unit are enforced on all branches in all test cases. All IEEE test cases are available in [39].		

TABLE II
TEST CASE RESULTS

Algorithm	Test Case	Max S^{mis} (MVA)	Obj. Val. Diff.	Min. Eigenvalue Ratio	Num. Iter.	Solver Time (sec)	Num. High-Order Buses
Algorithm 1: Iterative Solver	*case14Q*	1.08×10^{-3}	2.36×10^{-6}	1.08×10^{6}	3	36.4	(2^{nd}): 3, (3^{rd}): 0
	case14L	5.67×10^{-2}	2.84×10^{-6}	2.31×10^{4}	3	25.5	(2^{nd}): 4, (3^{rd}): 0
	case39Q	1.36×10^{-1}	1.01×10^{-4}	8.69×10^{3}	19	2857	(2^{nd}): 31, (3^{rd}): 2
	case39L	4.60×10^{-3}	8.12×10^{-7}	4.08×10^{5}	2	4.88	(2^{nd}): 2, (3^{rd}): 0
	case57Q	6.49×10^{-3}	4.35×10^{-6}	2.52×10^{5}	3	20.9	(2^{nd}): 4, (3^{rd}): 0
	case57L	8.76×10^{-4}	2.35×10^{-7}	2.18×10^{6}	2	88.8	(2^{nd}): 2, (3^{rd}): 0
	case118Q	2.13×10^{-1}	3.40×10^{-5}	6.78×10^{4}	3	172.6	(2^{nd}): 4, (3^{rd}): 0
	case118L	4.42×10^{-1}	4.59×10^{-5}	1.93×10^{4}	2	15.9	(2^{nd}): 2, (3^{rd}): 0
	case300	5.14×10^{-2}	3.74×10^{-6}	4.65×10^{4}	2	41.9	(2^{nd}): 2, (3^{rd}): 0
Locally Minimal Set of Higher-Order Buses	*case14Q*	7.64×10^{-3}	3.21×10^{-5}	1.52×10^{5}	N/A	10.1	(2^{nd}): 2, (3^{rd}): 0
	case14L	4.56×10^{-1}	4.80×10^{-4}	1.86×10^{3}	N/A	8.51	(2^{nd}): 2, (3^{rd}): 0
	case39Q	3.95×10^{-1}	7.42×10^{-4}	4.30×10^{3}	N/A	289.7	(2^{nd}): 22, (3^{rd}): 1
	case39L	1.13×10^{-2}	1.64×10^{-6}	1.85×10^{5}	N/A	3.21	(2^{nd}): 1, (3^{rd}): 0
	case57Q	1.02×10^{-2}	1.66×10^{-5}	1.14×10^{5}	N/A	7.67	(2^{nd}): 3, (3^{rd}): 0
	case57L	1.34×10^{-3}	6.60×10^{-7}	6.15×10^{5}	N/A	10.2	(2^{nd}): 1, (3^{rd}): 0
	case118Q	Same as Algorithm 1					
	case118L	Same as Algorithm 1					
	case300	Same as Algorithm 1					
First-Order Relaxation	*case14Q*	4.92×10^{0}	4.96×10^{-5}	2.70×10^{2}	N/A	0.69	N/A
	case14L	9.77×10^{0}	5.94×10^{-4}	1.41×10^{2}	N/A	0.87	N/A
	case39Q	1.34×10^{2}	3.54×10^{-2}	2.89×10^{2}	N/A	1.48	N/A
	case39L	6.91×10^{0}	5.55×10^{-6}	5.10×10^{3}	N/A	1.41	N/A
	case57Q	5.76×10^{0}	8.60×10^{-5}	1.78×10^{2}	N/A	2.08	N/A
	case57L	8.65×10^{0}	1.58×10^{-3}	3.84×10^{2}	N/A	3.53	N/A
	case118Q	7.36×10^{1}	2.16×10^{-4}	1.07×10^{2}	N/A	4.81	N/A
	case118L	1.07×10^{2}	7.53×10^{-3}	9.39×10^{1}	N/A	7.64	N/A
	case300	2.42×10^{1}	7.64×10^{-5}	1.29×10^{2}	N/A	17.9	N/A

problems. Thus, MATPOWER with the default interior point solver and default solver options finds the global optimum for all the test problems in this section, with optimal objective values listed in Table I and solution times of less than 0.5 seconds, but the first-order relaxation fails to solve each test problem. Second- and third-order moment relaxations *certify* that these solutions are the global optima.

Note that the existence of small examples for which both the first-order relaxation and traditional solution methods fail but the higher-order moment relaxations succeed in finding the global optimum (e.g., the five- and nine-bus systems in [6] and the five-bus system in [5]; see [20]–[22] for analysis of these and other small test cases) implies that similar phenomena can occur in large practical problems as well. This suggests the need for a wider variety of test problems, the development of which is beyond the scope of this paper.

Descriptions of each test problem are provided in Table I. Except for the IEEE 300-bus system, these problems are modifications of the IEEE test cases using one of two known methods for inducing failure of the first-order relaxation. Tightening apparent-power line-flow limits may induce failure of the first-order relaxation [16]. The letter "L" denotes corresponding problems. Decreasing the loading while reducing the generators' leading power factor range (i.e., decreasing the magnitude of the lower reactive power generation limits) may also induce failure of the first-order relaxation [6]. The letter "Q" denotes corresponding problems.

The results in this section are generated using a computer with a quad-core 2.70 GHz processor and 16 GB of RAM. The moment relaxations are implemented using MATLAB 2013a, YALMIP version 2014.02.21 [40], and Mosek version 7.0.0.102 [41].

Table II presents the results from applying the moment relaxations to the test problems in Table I. The first group of rows in Table II shows the results from Algorithm 1. The second group shows the results from an at-least-locally minimal set of higher-order buses. This set is derived by individually removing higher-order buses from the solution given by Algorithm 1. The third group shows the results from the first-order relaxation.

The columns show the values of three convergence metrics. The first metric is the maximum apparent-power injection

mismatch (Max S^{mis}), which has a 0.5 MVA tolerance. Note that the voltage magnitudes, power injections, and line flows satisfy the inequality constraints in (5) to within 0.005 per unit voltage and 0.5 MVA for both Algorithm 1 and the locally minimal set of higher-order buses.

The second metric compares the optimal objective value obtained directly from the moment relaxation, denoted as $C^{\text{Mom. Obj.}}$ and the cost implied by the voltage vector obtained from the closest matrix satisfying the rank condition, denoted as C^V. In the absence of numerical inaccuracy, these objective values are equal when the solution to the moment relaxation is exact. With imperfect solvers, this value may be non-zero even when the moment relaxation has an exact solution. The column Obj. Val. Diff. shows $\left|C^{\text{Mom. Obj.}} - C^V\right| / C^{\text{Mom. Obj.}}$, which is employed as the second convergence criterion with a tolerance of 1×10^{-3}. For a solution with small power injection mismatches, a small value in this column indicates global optimality for practical purposes.

Since the first-order relaxations yield large power injection mismatches, the objective value calculated from the voltage vector $\left(C^V\right)$ is not meaningful. Therefore, the Obj. Val. Diff. column for the first-order relaxation shows the difference between the globally optimal objective value $C^{\text{Mom. Obj.}}$ and the lower bound obtained from the objective value of the first-order relaxation C^{SDP} as $\left(C^{\text{Mom. Obj.}} - C^{\text{SDP}}\right) / C^{\text{Mom. Obj.}}$.

The third metric, which is first proposed in [28], is based directly on satisfaction of the rank condition. For a solution that satisfies the rank condition, the diagonal blocks of the moment matrices corresponding to the second-order terms have rank equal to one (i.e., the matrices $\mathbf{M}_1^{\mathcal{M}_m}(y)$ for $m = 1, \ldots, |\mathcal{M}|$ have a single non-zero eigenvalue). However, for numerical reasons, solvers do not yield a "hard zero" value for "zero" eigenvalues of these matrices. On the other hand, when the moment relaxation fails to be exact, these matrices have more than one non-zero eigenvalue.

To measure the satisfaction of the rank condition, we use the ratio between the largest and second-largest magnitude eigenvalues. The minimum such ratio among all the matrices $\mathbf{M}_1^{\mathcal{M}_m}(y)$ for $m = 1, \ldots, |\mathcal{M}|$, is termed the *minimum eigenvalue ratio*. If the solution to a case for which the relaxation was exact did have "hard zeros" for zero eigenvalues, the largest eigenvalue would be non-zero and the second-largest eigenvalue would be zero, resulting in a minimum eigenvalue ratio of infinity. In practice, numerical issues result in minimum eigenvalue ratios that are large (typical values are on the order of 10^4 for problems that satisfy the rank condition). Further, if the solution does not satisfy the rank condition, both the largest and second largest eigenvalues typically have similar magnitudes, therefore yielding a small value for the minimum eigenvalue ratio. Thus, a large value for the minimum eigenvalue ratio indicates satisfaction of the rank condition while a small value indicates failure to satisfy the rank condition.

Note that we do not use the minimum eigenvalue ratio as a convergence criteria. This is due to the fact that a solution to the moment relaxation can have a relatively poor (small) minimum eigenvalue ratio but the closest rank one matrix can still globally solve the OPF problem to within the other tolerances. (See, for instance, the results for *case14L* and *case39Q* in the second group of rows in Table II which show the results for the locally minimal set of higher-order buses.) We report the minimum eigenvalue ratio in the column Min. Eigenvalue Ratio of Table II.

Where applicable, Table II also shows the number of iterations of Algorithm 1 (Num. Iter.), the time spent in the MOSEK solver summed over all iterations (Solver Time), and the number of higher-order buses (Num. High-Order Bus), where the quantity in parentheses is the relaxation order.

Fig. 3 shows the lower bound on the objective function, the convergence metrics, and the solver times for each iteration of Algorithm 1. The lower bounds on the objective function in Fig. 3a are normalized so that the global optimum for each test problem has a value of one. Using a log scale, Fig. 3b shows the maximum power injection mismatch in MVA (i.e., the first convergence metric) for each iteration, Fig. 3c shows the objective value difference (i.e., the second convergence metric), and Fig. 3d shows the minimum eigenvalue ratio (i.e., the third convergence metric). Fig. 3e shows the solver time for each iteration of the algorithm.

We next emphasize several interesting observations from the results in Table II and Fig. 3. Solutions to the first-order relaxations generally have a subset of buses with large power injection mismatches, but the objective values are often quite close to the global optimum (see column Obj. Val. Diff. in Table II and the fact that the first iteration in Fig. 3a is approximately equal to 1 but the convergence metrics in Figs. 3b, 3c, and 3d are not satisfied). This suggests that there are often "hidden" or "nearly hidden" rank one solutions for the first-order relaxation [26]. Selective application of the higher-order constraints provides a mechanism for recovering these hidden rank one solutions. However, not all problems have hidden rank one solutions (e.g., the first-order relaxations of *case39Q* and *case118L*, which yield solutions that are 3.54% and 0.75%, respectively, below the global optimum, and the examples in [20]). The solver times from Algorithm 1 illustrate that additional computational effort is required to achieve the global solution relative to the lower bounds from the first-order relaxation.

The Num. Iter. column in Table II and the convergence metrics in Figs. 3b, 3c, and 3d show that only a small number of iterations are typically required to obtain a global optimum (i.e., the higher-order moment constraints are only required at a small number of buses). However, this is not always the case as *case39Q* requires many iterations and higher-order constraints at the majority of buses in the network. This problem demonstrates that the approach of selectively applying higher-order constraints to specific areas of the network may not be computationally tractable for all problems.

Since each iteration of Algorithm 1 adds constraints to the optimization problem, the cost shown in Fig. 3a should be non-decreasing. The cost is non-decreasing for all problems with the exception of iteration 17 for *case39Q*. The decrease at iteration 17 is explained by the fact that the semidefinite programming solver does not converge to a sufficient tolerance (i.e., the constraints are not satisfied) at this iteration of Algorithm 1.

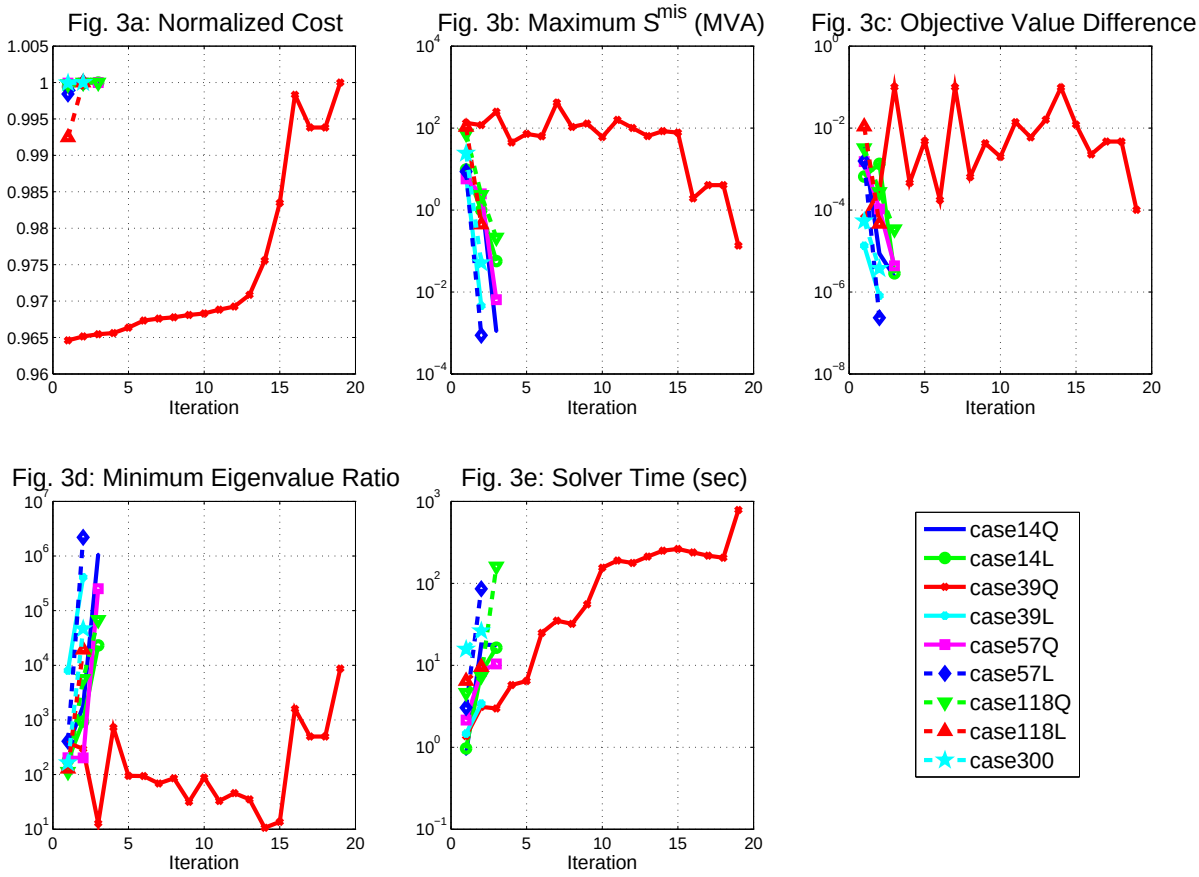

Fig. 3. Test Case Results for Each Iteration of Algorithm 1. Fig. 3a shows the lower bound on the optimal objective value normalized by the globally optimal objective value. Figs. 3b, 3c, and 3d show the three convergence metrics, and Fig. 3e shows the solver time for each iteration.

For *case118Q*, *case118L*, and *case300*, Algorithm 1 finds an at-least-locally minimal set of buses requiring the higher-order relaxation. For *case14Q*, *case39L*, *case57Q*, and *case57L*, Algorithm 1 only uses one more bus than a locally minimal set. While this indicates that the heuristic with $h = 2$ is effective in identifying a minimal or near-minimal set of higher-order buses for many OPF problems, we emphasize that these are only known to be locally minimal; a different heuristic may identify a smaller set of buses. Further, Algorithm 1 does not identify a near-minimal set of higher-order buses for *case39Q*, which results in an almost order-of-magnitude larger solution time than necessary. More sophisticated heuristics could lead to better performance for some problems.

Finally, note that solution times have stronger dependence on the number of higher-order buses, the size of the maximal cliques they are contained within, and the relaxation orders required than the size of the system. For instance, although *case300* has 7.7 times more buses, Algorithm 1 has a factor of 68.2 greater solution time for *case39Q* due to the number of buses with second- and third-order constraints in *case39Q*. Further note that although solution times for the moment relaxations are not yet competitive with mature local solvers, such as the interior point method in MATPOWER which solved all test problems in Table I in less than 0.5 seconds, the moment relaxations with Algorithm 1 provide a computationally tractable approach for globally solving many problems for which the first-order relaxation fails to yield a global solution. The moment relaxations also *certify* global optimality in contrast to traditional solvers which only guarantee a local optimum.

VI. Conclusion

While existing convex relaxations globally solve many OPF problems, there are practical problems for which existing relaxations fail to yield physically meaningful solutions. This paper has described a hierarchy of "moment" relaxations that globally solve many problems for which existing relaxations fail. The moment relaxations, which take the form of semidefinite programs, are developed from the Lasserre hierarchy for generalized moment problems. Increasing the order in this hierarchy results in "tighter" relaxations at the computational cost of larger semidefinite programs.

Solving the moment relaxations for larger problems requires both exploiting power system sparsity and selectively applying the higher-order moment relaxation constraints. A matrix completion decomposition for exploiting sparsity was

first presented. Next, taking advantage of the observation that first-order relaxations are sufficient for large regions of typical OPF problems, this paper proposed an iterative algorithm for solving the moment relaxations. A heuristic at each iteration identifies where to enforce the higher-order relaxation constraints. The proposed algorithm's effectiveness was demonstrated by globally solving several test cases for which existing convex relaxations failed.

Future work includes improving the computational performance of the proposed algorithm. Alternative heuristics for determining where to apply the higher-order constraints may reduce solution times. Distributed solution algorithms, which have proven valuable for existing relaxations [42], may also speed computation of the moment relaxations.

Identification and exploration of realistic test cases for which low-order relaxations fail is another important future research direction. The problems in [15] used to demonstrate that the OPF problem is, in general, NP-hard serve as example cases for which low-order moment relaxations fail to yield a global solution [20]. However, with a large number of global optima, these problems are very atypical. Extending the work of, e.g., [23]–[25], sufficient conditions for tightness of the moment relaxations would also be valuable contributions.

Additional future work also includes application of moment relaxations to more general OPF formulations that include, for instance, discrete devices, security constraints, and transient stability constraints. Exploiting synergies with robust and chance-constrained optimization techniques for the OPF problem appears particularly promising [43]. Extension to other problems in power system engineering, such as state estimation, voltage stability margins, and power flow calculations, is another avenue of future work.

ACKNOWLEDGMENT

The authors acknowledge the support of the Dow Postdoctoral Fellowship in Sustainability, ARPA-E grant DE-AR0000232, and Los Alamos National Laboratory subcontract 270958.

REFERENCES

[1] F. Capitanescu, J. M. Ramos, P. Panciatici, D. Kirschen, A. M. Marcolini, L. Platbrood, and L. Wehenkel, "State-of-the-Art, Challenges, and Future Trends in Security Constrained Optimal Power Flow," *Electric Power Systems Research*, vol. 81, no. 8, pp. 1731 – 1741, 2011.

[2] B. Stott and O. Alsaç, "Optimal Power Flow–Basic Requirements for Real-Life Problems and their Solutions," in *SEPOPE XII Symposium, Rio de Janeiro, Brazil*, 2012.

[3] D. Phan and J. Kalagnanam, "Some Efficient Optimization Methods for Solving the Security-Constrained Optimal Power Flow Problem," *IEEE Trans. Power Syst.*, vol. 29, no. 2, pp. 863–872, Mar. 2014.

[4] L. Platbrood, F. Capitanescu, C. Merckx, H. Crisciu, and L. Wehenkel, "A Generic Approach for Solving Nonlinear-Discrete Security-Constrained Optimal Power Flow Problems in Large-Scale Systems," *IEEE Trans. Power Syst.*, vol. 29, no. 3, pp. 1194–1203, May 2014.

[5] B. Lesieutre and I. Hiskens, "Convexity of the Set of Feasible Injections and Revenue Adequacy in FTR Markets," *IEEE Trans. Power Syst.*, vol. 20, no. 4, pp. 1790–1798, Nov. 2005.

[6] W. Bukhsh, A. Grothey, K. McKinnon, and P. Trodden, "Local Solutions of the Optimal Power Flow Problem," *IEEE Trans. Power Syst.*, vol. 28, no. 4, pp. 4780–4788, 2013.

[7] M. Huneault and F. Galiana, "A Survey of the Optimal Power Flow Literature," *IEEE Trans. Power Syst.*, vol. 6, no. 2, pp. 762–770, May 1991.

[8] J. Momoh, R. Adapa, and M. El-Hawary, "A Review of Selected Optimal Power Flow Literature to 1993. I. Nonlinear and Quadratic Programming Approaches," *IEEE Trans. Power Syst.*, vol. 14, no. 1, pp. 96–104, Feb. 1999.

[9] J. Momoh, M. El-Hawary, and R. Adapa, "A Review of Selected Optimal Power Flow Literature to 1993. II. Newton, Linear Programming and Interior Point Methods," *IEEE Trans. Power Syst.*, vol. 14, no. 1, pp. 105–111, Feb. 1999.

[10] R. Zimmerman, C. Murillo-Sánchez, and R. Thomas, "MATPOWER: Steady-State Operations, Planning, and Analysis Tools for Power Systems Research and Education," *IEEE Trans. Power Syst.*, no. 99, pp. 1–8, 2011.

[11] S. Frank, I. Steponavice, and S. Rebennack, "Optimal Power Flow: A Bibliographic Survey, Parts I and II," *Energy Syst.*, vol. 3, no. 3, pp. 221–289, 2012.

[12] A. Castillo and R. O'Neill, "Survey of Approaches to Solving the ACOPF (OPF Paper 4)," US Federal Energy Regulatory Commission, Tech. Rep., Mar. 2013.

[13] D. Molzahn, B. Lesieutre, and C. DeMarco, "A Sufficient Condition for Global Optimality of Solutions to the Optimal Power Flow Problem," *IEEE Trans. Power Syst.*, vol. 29, no. 2, pp. 978–979, Mar. 2014.

[14] A. Castillo and R. O'Neill, "Computational Performance of Solution Techniques Applied to the ACOPF (OPF Paper 5)," US Federal Energy Regulatory Commission, Tech. Rep., Jan. 2013.

[15] J. Lavaei and S. Low, "Zero Duality Gap in Optimal Power Flow Problem," *IEEE Trans. Power Syst.*, vol. 27, no. 1, pp. 92–107, Feb. 2012.

[16] D. K. Molzahn, B. C. Lesieutre, and C. L. DeMarco, "Investigation of Non-Zero Duality Gap Solutions to a Semidefinite Relaxation of the Power Flow Equations," in *47th Hawaii Int. Conf. Syst. Sci. (HICSS)*, 6-9 Jan. 2014.

[17] D. Molzahn, B. Lesieutre, and C. DeMarco, "A Sufficient Condition for Power Flow Insolvability With Applications to Voltage Stability Margins," *IEEE Trans. Power Syst.*, vol. 28, no. 3, pp. 2592–2601, 2013.

[18] J. B. Lasserre, "Global Optimization with Polynomials and the Problem of Moments," *SIAM J. on Optimiz.*, vol. 11, no. 3, pp. 796–817, 2001.

[19] J.-B. Lasserre, *Moments, Positive Polynomials and Their Applications*. Imperial College Press, 2010, vol. 1.

[20] D. K. Molzahn and I. A. Hiskens, "Moment-Based Relaxation of the Optimal Power Flow Problem," *18th Power Syst. Comput. Conf. (PSCC)*, 18-22 Aug. 2014, preprint available: http://arxiv.org/abs/1312.1992.

[21] C. Josz, J. Maeght, P. Panciatici, and J. C. Gilbert, "Application of the Moment-SOS Approach to Global Optimization of the OPF Problem," To appear in *IEEE Trans. Power Syst.*, preprint available: http://arxiv.org/abs/1311.6370.

[22] B. Ghaddar, J. Marecek, and M. Mevissen, "Optimal Power Flow as a Polynomial Optimization Problem," IBM Research Ireland, Tech. Rep., 2014, http://arxiv.org/abs/1404.3626.

[23] S. H. Low, "Convex Relaxation of Optimal Power Flow–Parts I: Formulations and Equivalence," *IEEE Trans. Control Network Syst.*, vol. 1, no. 1, pp. 15–27, Mar. 2014.

[24] S. Low, "Convex Relaxation of Optimal Power Flow–Part II: Exactness," *IEEE Trans. Control Network Syst.*, vol. 1, no. 2, pp. 177–189, Jun. 2014.

[25] R. Madani, S. Sojoudi, and J. Lavaei, "Convex Relaxation for Optimal Power Flow Problem: Mesh Networks," To appear in *IEEE Trans. Power Syst.*, 2014.

[26] R. Louca, P. Seiler, and E. Bitar, "A Rank Minimization Algorithm to Enhance Semidefinite Relaxations of Optimal Power Flow," in *51st Annu. Allerton Conf. Commun., Control, and Comput.*, 2-4 Oct. 2013.

[27] R. Jabr, "Exploiting Sparsity in SDP Relaxations of the OPF Problem," *IEEE Trans. Power Syst.*, vol. 27, no. 2, pp. 1138–1139, May 2012.

[28] D. Molzahn, J. Holzer, B. Lesieutre, and C. DeMarco, "Implementation of a Large-Scale Optimal Power Flow Solver Based on Semidefinite Programming," *IEEE Trans. Power Syst.*, vol. 28, no. 4, pp. 3987–3998, 2013.

[29] H. Waki, S. Kim, M. Kojima, and M. Muramatsu, "Sums of Squares and Semidefinite Program Relaxations for Polynomial Optimization Problems with Structured Sparsity," *SIAM J. Optimiz.*, vol. 17, no. 1, pp. 218–242, 2006.

[30] J. Lavaei, "Zero Duality Gap for Classical OPF Problem Convexifies Fundamental Nonlinear Power Problems," in *American Control Conference (ACC), 2011*, June 2011, pp. 4566–4573.

[31] M. Vrakopoulou, M. Katsampani, K. Margellos, J. Lygeros, and G. Andersson, "Probabilistic Security-Constrained AC Optimal Power Flow," in *PowerTech (POWERTECH), 2013 IEEE Grenoble*, June 2013, pp. 1–6.

[32] M. Baradar, M. Hesamzadeh, and M. Ghandhari, "Second-Order Cone Programming for Optimal Power Flow in VSC-Type AC-DC Grids," *IEEE Trans. Power Syst.*, vol. 28, no. 4, pp. 4282–4291, Nov. 2013.
[33] W. A. Bukhsh, A. Grothey, K. I. McKinnon, and P. A. Trodden, "Local Solutions of Optimal Power Flow," University of Edinburgh School of Mathematics, Tech. Rep. ERGO 11-017, 2011, [Online]. Available: http://www.maths.ed.ac.uk/ERGO/pubs/ERGO-11-017.html.
[34] M. Andersen, A. Hansson, and L. Vandenberghe, "Reduced-Complexity Semidefinite Relaxations of Optimal Power Flow Problems," *IEEE Trans. Power Syst.*, vol. PP, no. 99, pp. 1–9, 2014.
[35] R. Gron, C. Johnson, E. Sá, and H. Wolkowicz, "Positive Definite Completions of Partial Hermitian Matrices," *Linear Algebra Appl.*, vol. 58, pp. 109–124, 1984.
[36] M. Fukuda, M. Kojima, K. Murota, K. Nakata *et al.*, "Exploiting Sparsity in Semidefinite Programming via Matrix Completion I: General Framework," *SIAM J. Optimiz.*, vol. 11, no. 3, pp. 647–674, 2001.
[37] R. Tarjan and M. Yannakakis, "Simple Linear-Time Algorithms to Test Chordality of Graphs, Test Acyclicity of Hypergraphs, and Selectively Reduce Acyclic Hypergraphs," *SIAM J. Comput.*, vol. 13, p. 566, 1984.
[38] T. Davis, J. Gilbert, S. Larimore, and E. Ng, "Algorithm 836: COLAMD, a Column Approximate Minimum Degree Ordering Algorithm," *ACM T. Math. Software*, vol. 30, no. 3, pp. 377–380, Sept. 2004.
[39] Power Systems Test Case Archive. University of Washington Department of Electrical Engineering. [Online]. Available: http://www.ee.washington.edu/research/pstca/
[40] J. Lofberg, "YALMIP: A Toolbox for Modeling and Optimization in MATLAB," in *IEEE Int. Symp. Compu. Aided Control Syst. Des.*, 2004, pp. 284–289.
[41] The MOSEK Optimization Toolbox for MATLAB Manual version 7.0. MOSEK ApS. [Online]. Available: http://docs.mosek.com/7.0/toolbox/
[42] A. Lam, B. Zhang, and D. Tse, "Distributed Algorithms for Optimal Power Flow Problem," in *51st Annu. Conf. Decis. Control (CDC)*, 10-13 Dec. 2012.
[43] P. Panciatici, M. Campi, S. Garatti, S. Low, D. Molzahn, A. Sun, and L. Wehenkel, "Advanced optimization methods for power systems," *18th Power Syst. Comput. Conf. (PSCC)*, 18-22 Aug. 2014.

Daniel K. Molzahn (S'09-M'13) is a Dow Sustainability Fellow at the University of Michigan, Ann Arbor. He received the B.S., M.S., and Ph.D. degrees in electrical engineering and the Masters of Public Affairs degree from the University of Wisconsin-Madison, where he was a National Science Foundation Graduate Research Fellow. His research interests are in the application of optimization techniques and policy analysis to electric power systems.

Ian A. Hiskens (S'77-M'80-SM'96-F'06) received the B.Eng. degree in electrical engineering and the B.App.Sc. degree in mathematics from the Capricornia Institute of Advanced Education, Rockhampton, Australia, in 1980 and 1983 respectively, and the Ph.D. degree in electrical engineering from the University of Newcastle, Australia, in 1991.

He is the Vennema Professor of Engineering in the Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor. He has held prior appointments in the Queensland electricity supply industry, and various universities in Australia and the United States. His research interests lie at the intersection of power system analysis and systems theory, with recent activity focused largely on integration of renewable generation and controllable loads.

Dr. Hiskens is actively involved in various IEEE societies, and is VP-Finance of the IEEE Systems Council. He is a Fellow of the IEEE, a Fellow of Engineers Australia, and a Chartered Professional Engineer in Australia.